AMENDMENT AND RESPONSE UNDER 37 CFR § 1.111
Serial Number: 10/643,586
Filing Date: August 18, 2003
Title: DECOUPLED SCALAR/VECTOR COMPUTER ARCHITECTURE SYSTEM AND METHOD (As Amended)

Page 2
Dkt: 1376.699US1

IN THE CLAIMS

Please amend the claims as follows:

1.   (Currently Amended) In a computer system having a scalar processing unit and a vector processing unit, wherein the vector processing unit includes a vector dispatch unit and wherein the vector dispatch unit includes a predispatch queue and a dispatch queue, a method of decoupling operation of the scalar processing unit from that of the vector processing unit, the method comprising:

dispatching ~~a~~ vector ~~instruction~~ instructions from the scalar processing unit to the vector dispatch unit, wherein dispatching includes sending the vector ~~instruction~~ instructions from the scalar processing unit to the vector dispatch unit even if all scalar operands are not ready and even if all scalar instructions issued prior to the vector instructions are not scalar committed;

queueing up the vector instructions received from the scalar processing unit in the vector dispatch unit's predispatch queue;

reading scalar operands from the scalar processing unit, wherein reading includes transferring the scalar operands from the scalar processing unit to the vector dispatch unit;

predispatching~~, within the vector dispatch unit,~~ the vector ~~instruction received from the scalar processing unit if all previously received vector~~ instructions from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed;

dispatching the predispatched vector instruction from the ~~vector dispatch unit~~ dispatch queue if all required scalar operands are ready; and

executing the vector instruction dispatched from the ~~vector dispatch unit~~ dispatch queue as a function of the scalar operands.

2.   (Previously Presented) The method according to claim 1, wherein executing the dispatched vector instruction includes translating an address associated with the vector instruction and trapping on a translation fault.

3.   (Currently Amended) In a computer system having a scalar processing unit and a vector processing unit, wherein the vector processing unit includes a vector dispatch unit and wherein the vector dispatch unit includes a predispatch queue and a dispatch queue, a method of decoupling operation of the scalar processing unit from that of the vector processing unit, the method comprising:

     dispatching a vector ~~instruction~~ instructions from the scalar processing unit to the vector dispatch unit, wherein dispatching includes sending the vector ~~instruction~~ instructions from the scalar processing unit to the vector dispatch unit even if all scalar operands are not ready and even if all scalar instructions issued prior to the vector instructions are not scalar committed;

     queueing up the vector instructions received from the scalar processing unit in the vector dispatch unit's predispatch queue;

     reading scalar operands from the scalar processing unit, wherein reading includes transferring the scalar operands from the scalar processing unit to the vector dispatch unit;

     predispatching~~, within the vector dispatch unit,~~ the vector ~~instruction received from the scalar processing unit if all previously received vector~~ instructions from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed;

     dispatching the predispatched vector instruction from the ~~vector dispatch unit~~ dispatch queue if all required scalar operands are ready;

     generating an address for a vector load;

     issuing a vector load request to memory;

     receiving vector data from memory;

     storing the vector data in a load buffer;

AMENDMENT AND RESPONSE UNDER 37 CFR § 1.111
Serial Number: 10/643,586
Filing Date: August 18, 2003
Title: DECOUPLED SCALAR/VECTOR COMPUTER ARCHITECTURE SYSTEM AND METHOD (As Amended)

Page 4
Dkt: 1376.699US1

transferring the vector data from the load buffer to a vector register; and

executing the vector instruction dispatched from the vector ~~vector dispatch unit~~ <u>dispatch queue</u> on the vector data stored in the vector register.

4. (Original) The method according to claim 3, wherein the vector processing unit includes a vector execute unit and a vector load/store unit, wherein issuing a vector load request to memory includes issuing and executing vector memory references in the vector load/store unit when the vector load store unit has received the instruction and memory operands from the scalar processing unit.

5. (Currently Amended) In a computer system having a scalar processing unit and a vector processing unit, wherein the vector processing unit includes a vector dispatch unit <u>and wherein the vector dispatch unit includes a predispatch queue and a dispatch queue</u>, a method of decoupling operation of the scalar processing unit from that of the vector processing unit, the method comprising:

dispatching ~~a~~ vector ~~instruction~~ <u>instructions</u> from the scalar processing unit to the vector dispatch unit, wherein dispatching includes sending the vector ~~instruction~~ <u>instructions</u> from the scalar processing unit to the vector dispatch unit even if all scalar operands are not ready <u>and even if all scalar instructions issued prior to the vector instructions are not scalar committed</u>;

<u>queueing up the vector instructions received from the scalar processing unit in the vector dispatch unit's predispatch queue;</u>

reading scalar operands from the scalar processing unit, wherein reading includes transferring the scalar operands from the scalar processing unit to the vector dispatch unit;

predispatching~~, within the vector dispatch unit,~~ the vector ~~instruction received from the scalar processing unit if all previously received vector~~ instructions <u>from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed;</u>

dispatching the predispatched vector instruction from the ~~vector dispatch unit~~ dispatch queue if all required scalar operands are ready;

generating a first and a second address for a vector load;

issuing first and second vector load requests to memory;

receiving vector data associated with the first and second addresses from memory;

storing vector data associated with the first address in a first vector register;

storing vector data associated with the second address in a second vector register;

executing a vector instruction on the vector data stored in the first vector register;

renaming the second vector register; and

executing the vector instruction dispatched from the ~~vector dispatch unit~~ dispatch queue on the vector data stored in the renamed vector register.

6.   (Previously Presented) The method according to claim 5, wherein the vector processing unit includes a vector execute unit and a vector load/store unit, wherein issuing a vector load request to memory includes issuing and executing vector memory references in the vector load/store unit when the vector load store unit has received the instruction and memory operands from the scalar processing unit.

7.   (Currently Amended) A computer system, comprising:

a scalar processing unit; and

a vector processing unit, wherein the vector processing unit includes a vector dispatch unit, a vector execute unit and a vector load/store unit, wherein the vector dispatch unit includes a predispatch queue and a dispatch queue;

wherein the scalar processing unit dispatches ~~a vector~~ instruction instructions to the vector dispatch unit even if all scalar operands associated with the instructions are not ready;

wherein the scalar processing unit reads scalar operands and transfers the read scalar operands from the scalar processing unit to the vector dispatch unit;

wherein the vector dispatch unit stores vector instructions received from the scalar processing unit in the vector dispatch unit's predispatch queue, predispatches ~~, within the vector dispatch unit,~~ the vector ~~instruction received from the scalar processing~~ if all previously

~~received vector~~ instructions <u>from the predispatch queue to the dispatch queue in the order</u>
<u>received;</u>

<u>wherein the vector dispatch unit determines if all scalar instructions issued prior to the</u>
<u>vector instruction at the head of the predispatch queue</u> are scalar committed and <u>transfers the</u>
<u>vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions</u>
<u>issued prior to the vector instruction at the head of the predispatch queue are scalar committed;</u>

<u>wherein the vector dispatch unit</u> ~~then~~ dispatches the predispatched vector instruction from
the vector dispatch unit to one or more of the vector execute unit and the vector load/store unit if
all required scalar operands are ready;

wherein the vector load/store unit receives an instruction and memory operands from the
scalar processing unit, issues and executes a vector memory load reference as a function of the
instruction and the memory operands received from the scalar processing unit, and stores data
received as a result of the vector memory reference in a load buffer; and

wherein the vector execute unit issues the vector memory load instruction and transfers
the data received as a result of the vector memory reference from the load buffer to a vector
register.

8.   (Currently Amended) In a computer system having a scalar processing unit and a vector
processing unit<u>, wherein the vector processing unit includes a vector dispatch unit and wherein</u>
<u>the vector dispatch unit includes a predispatch queue and a dispatch queue</u>, a method of
decoupling scalar and vector execution, comprising:

dispatching scalar instructions to a scalar instruction queue in the scalar processing unit;

dispatching a vector instruction that requires scalar operands from the scalar processing
unit to the scalar instruction queue and to ~~a vector instruction~~ <u>the predispatch</u> queue, wherein
dispatching includes sending the vector instruction from scalar processing unit to the ~~vector~~
~~dispatch unit~~ <u>predispatch queue</u> even if all scalar operands are not ready;

<u>queueing up two or more vector instructions received from the scalar processing unit in</u>
<u>the predispatch queue;</u>

executing the vector instruction in the scalar processing unit, wherein executing the vector instruction in the scalar processing unit includes writing a scalar operand to a scalar operand queue;

predispatching, within the vector dispatch unit, the vector ~~instruction received from the scalar processing unit if all previously received vector~~ instructions <u>from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue</u> are scalar committed <u>and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed</u>;

notifying the vector processing unit that the scalar operand is available in the scalar operand queue;

dispatching the predispatched vector instruction from <u>the dispatch queue of</u> the vector dispatch unit, if all required scalar operands are ready; and

executing the vector instruction dispatched from the vector dispatch unit in the vector processing unit, wherein executing the vector instruction in the vector processing unit includes reading the scalar operand from the scalar operand queue.

9.   (Currently Amended) In a computer system having a scalar processing unit and a vector processing unit<u>, wherein the vector processing unit includes a vector dispatch unit and wherein the vector dispatch unit includes a predispatch queue and a dispatch queue</u>, a method of decoupling a vector memory reference and a vector execution, comprising:

dispatching scalar instructions to a scalar instruction queue in the scalar processing unit;

dispatching a vector instruction from the scalar processing unit to the scalar instruction queue and to ~~a vector instruction~~ <u>the predispatch</u> queue in the vector processing unit, wherein dispatching includes sending the vector instruction from the scalar processing unit to the ~~vector instruction~~ <u>predispatch</u> queue in the vector processing unit even if all scalar operands are not ready;

<u>queueing up two or more vector instructions received from the scalar processing unit in the predispatch queue;</u>

AMENDMENT AND RESPONSE UNDER 37 CFR § 1.111
Serial Number: 10/643,586
Filing Date: August 18, 2003
Title: DECOUPLED SCALAR/VECTOR COMPUTER ARCHITECTURE SYSTEM AND METHOD (As Amended)

Page 8
Dkt: 1376.699US1

executing the vector instruction in the scalar processing unit, wherein executing the vector instruction in the scalar processing unit includes generating an address and writing the address to a scalar operand queue;

predispatching, within the vector dispatch unit, the vector ~~instruction sent received from the scalar processing unit if all previously received vector~~ instructions from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed;

notifying the vector processing unit that the address is available in the scalar operand queue;

dispatching the predispatched vector instruction from the dispatch queue of the vector processing unit if all required scalar operands are ready; and

executing the vector instruction dispatched from the vector processing unit in the vector processing unit, wherein executing the vector instruction in the vector processing unit includes reading the address from the scalar operand queue and generating a memory request as a function of the address read from the scalar operand queue.

10.   (Currently Amended) In a computer system having a scalar processing unit and a vector processing unit, wherein the vector processing unit includes a vector dispatch unit and wherein the vector dispatch unit includes a predispatch queue and a dispatch queue, a method of executing a vector instruction, comprising:

dispatching scalar instructions to a scalar instruction queue in the scalar processing unit;

dispatching a vector instruction from the scalar processing unit to the scalar instruction queue and to ~~a vector instruction~~ the predispatch queue in the vector processing unit, wherein dispatching includes sending the vector instruction from the scalar processing unit to the ~~vector instruction~~ predispatch queue in the vector processing unit even if all scalar operands are not ready;

queueing up two or more vector instructions received from the scalar processing unit in the predispatch queue;

executing the vector instruction in the scalar processing unit, wherein executing the vector instruction in the scalar processing unit includes generating an address and writing the address to a scalar operand queue;

predispatching, within the vector dispatch unit, the vector ~~instruction sent received from the scalar processing unit if all previously received vector~~ instructions from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed;

notifying the vector processing unit that the address is available in the scalar operand queue;

dispatching the predispatched vector instruction from the dispatch queue of the vector processing unit if all required scalar operands are ready; and

executing the vector instruction dispatched from the vector processing unit in the vector processing unit, wherein executing the dispatched vector instruction in the vector processing unit includes:

reading the address from the scalar operand queue;

generating a memory request as a function of the address read from the scalar operand queue;

receiving vector data from memory;

storing the vector data in a load buffer;

transferring the vector data from the load buffer to a vector register; and

executing a vector instruction on the vector data stored in the vector register.

11.   (Currently Amended) In a computer system having a scalar processing unit and a vector processing unit, wherein the vector processing unit includes a vector dispatch unit and wherein

the vector dispatch unit includes a predispatch queue and a dispatch queue, a method of unrolling a loop, comprising:

preparing a first and second vector instruction, wherein each vector instruction ~~execute~~ executes an iteration through the loop and wherein each vector instruction requires calculation of a scalar loop value;

dispatching the first and second vector instructions from the scalar processing unit to a scalar instruction queue in the scalar processing unit and to ~~a vector instruction~~ the predispatch queue in the vector processing unit, wherein dispatching includes sending the first and second vector instructions from the scalar processing unit to the predispatch queue in the vector processing unit even if all scalar operands are not ready;

queueing up the first and second vector instructions received from the scalar processing unit in the predispatch queue;

executing a portion of each vector instruction in the scalar processing unit, wherein executing a portion of each vector instruction in the scalar processing unit includes writing a scalar operand representing the scalar loop value calculated for each vector instruction to a scalar operand queue;

predispatching, within the vector processing unit, each vector instruction received from the scalar processing unit if all previously received vector instructions are scalar committed;

predispatching, within the vector dispatch unit, each vector instruction ~~sent received from the scalar processing unit if all previously received vector instructions~~ from the predispatch queue to the dispatch queue in the order received, wherein predispatching includes determining if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed and transferring the vector instruction from the predispatch queue to the dispatch queue only if all scalar instructions issued prior to the vector instruction at the head of the predispatch queue are scalar committed;

notifying the vector processing unit that the scalar operand is available in the scalar operand queue;

dispatching the predispatched vector instruction from the dispatch queue of the vector processing unit if all required scalar operands are ready; and

AMENDMENT AND RESPONSE UNDER 37 CFR § 1.111
Serial Number: 10/643,586
Filing Date: August 18, 2003
Title: DECOUPLED SCALAR/VECTOR COMPUTER ARCHITECTURE SYSTEM AND METHOD (As Amended)

Page 11
Dkt: 1376.699US1

executing the first and second vector instructions dispatched from the ~~vector processing unit in~~ dispatch queue of the vector processing unit, wherein executing the dispatched vector instruction in the vector processing unit includes reading the scalar operands associated with each instruction from the scalar operand queue.

12.   (Previously Presented) The method according to claim 3, wherein storing the vector data in a load buffer and transferring the vector data from the load buffer to a vector register are decoupled from each other.

13.   (Previously Presented) The method according to claim 3, wherein storing the vector data in a load buffer includes writing memory load data to the load buffer until all previous memory operations complete without fault.

14.   (Previously Presented) The method according to claim 4, wherein storing the vector data in a load buffer and transferring the vector data from the load buffer to a vector register are decoupled from each other.

15.   (Previously Presented) The method according to claim 4, wherein storing the vector data in a load buffer includes writing memory load data to the load buffer until all previous memory operations complete without fault.

16.   (Previously Presented) The system according to claim 7, wherein the load buffer stores memory load data until it is determined that no previous memory operation will fail and, if no previous memory operations have failed, the load buffer transfers the data to the vector register.

17.   (Currently Amended) The method according to claim 1, wherein the method further includes [:] marking the vector instruction as complete, wherein marking the vector instruction as complete includes:

AMENDMENT AND RESPONSE UNDER 37 CFR § 1.111
Serial Number: 10/643,586
Filing Date: August 18, 2003
Title: DECOUPLED SCALAR/VECTOR COMPUTER ARCHITECTURE SYSTEM AND METHOD (As Amended)

Page 12
Dkt: 1376.699US1

if the vector instruction is not a memory instruction and if the vector instruction does not require scalar operands, indicating that the vector instruction is complete when the vector instruction is dispatched from the scalar processing unit;

if the vector instruction is not a memory instruction but requires scalar operands, indicating that the vector instruction is complete when the scalar operands are available; and

if the vector instruction is a memory instruction, indicating that the vector instruction is complete when the vector address has been translated; and

graduating the vector instruction if the vector instruction is marked complete.

18.   (Canceled)

19.   (Currently Amended) The method according to claim 8, wherein the method further includes [:] marking the vector instruction as complete, wherein marking the vector instruction as complete includes:

if the vector instruction is not a memory instruction and if the vector instruction does not require scalar operands, indicating that the vector instruction is complete when the vector instruction is dispatched from the scalar processing unit;

if the vector instruction is not a memory instruction but requires scalar operands, indicating that the vector instruction is complete when the scalar operands are available; and

if the vector instruction is a memory instruction, indicating that the vector instruction is complete when the vector address has been translated; and

graduating the vector instruction if the vector instruction is marked complete.

20.   (Currently Amended) The method according to claim 9, wherein the method further includes [:] marking the vector instruction as complete, wherein marking the vector instruction as complete includes:

if the vector instruction is not a memory instruction and if the vector instruction does not require scalar operands, indicating that the vector instruction is complete when the vector instruction is dispatched from the scalar processing unit;

if the vector instruction is not a memory instruction but requires scalar operands, indicating that the vector instruction is complete when the scalar operands are available; and

if the vector instruction is a memory instruction, indicating that the vector instruction is complete when the vector address has been translated; and

graduating the vector instruction if the vector instruction is marked complete.